arXiv:1609.06782v1 [cs.CV] 22 Sep 2016

Deep Learning for Video Classification and Captioning

Zuxuan Wu[1,3], Ting Yao[2], Yanwei Fu[1], Yu-Gang Jiang[1]

[1]*Fudan University,* [2]*Microsoft Research Asia,* [3]*University of Maryland*
zxwu@cs.umd.edu, tiyao@microsoft.com, {ygj, yanweifu}@fudan.edu.cn

1. Introduction

Today's digital contents are inherently multimedia: text, audio, image, video and *etc.* Video, in particular, becomes a new way of communication between Internet users with the proliferation of sensor-rich mobile devices. Accelerated by the tremendous increase in Internet bandwidth and storage space, video data has been generated, published and spread explosively, becoming an indispensable part of today's big data. This has encouraged the development of advanced techniques for a broad range of video understanding applications. A fundamental issue that underlies the success of these technological advances is the understanding of video contents. Recent advances in deep learning in image [41, 68, 17, 50] and speech [21, 27] domain have encouraged techniques to learn robust video feature representations to effectively exploit abundant multimodal clues in video data.

In this paper, we focus on reviewing two lines of research aiming to stimulate the comprehension of videos with deep learning: video classification and video captioning. While video classification concentrates on automatically labeling video clips based on their semantic contents like human actions or complex events, video captioning attempts to generate a complete and natural sentence, enriching the single label as in video classification, to capture the most informative dynamics in videos.

There have been several efforts surveying literatures on video content understanding. Most of the approaches surveyed in these works adopted hand-crafted features coupled with typical machine learning pipelines for action recognition and event detection [1, 88, 61, 35]. In contrast, this paper focuses on discussing state-of-the-art deep learning techniques not only for video classification but also video captioning. As deep learning for video analysis is an emerging and vibrant field, we hope this paper could help stimulate future research along the line.

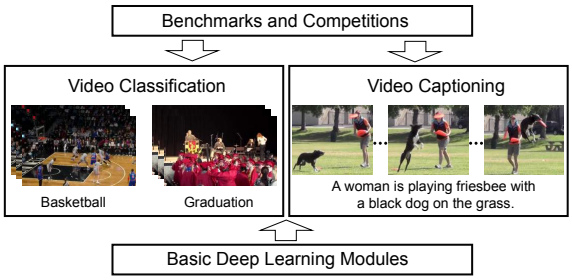

Figure 1: An overview of the organization of this paper.

Figure 1 shows the organization of this paper. To make the paper self-contained, we first introduce the basic modules that are widely adopted in state-of-the-art deep learning pipelines in Sec. 2. After that, we discuss representative works on video classification and video captioning in Sec. 3 and Sec. 4 respectively. Finally, we provide a review of popular benchmarks and competitions in Sec. 5, which are critical for evaluating the technical progress of this vibrant field.

2. Basic Deep Learning Modules

In this section, we briefly review basic deep learning modules that have been widely adopted in the literatures for video analysis.

2.1. Convolutional Neural Networks (CNNs)

Inspired by the visual perception mechanisms of animals [30] and McCulloch-Pitts Model [53], Fukushima proposed the "neocognitron" in 1980, which is the first computational model of using local connectivities between neurons of a hierarchically transformed image [16]. To obtain the translational invariance, Fukushima applied neurons with the same parameters on patches of the previous layer at different locations; thus this can be considered as the predecessor of CNN. Further inspired by this idea, LeCun *et al.* designed and trained the modern framework of CNN – LeNet-5 and obtained the state-of-the-art performance on several pattern recognition datasets, *e.g.*, handwritten character recognition [45]. LeNet-5 has multiple layers and is trained with the back-propagation algorithm in an end-to-end formulation, *i.e.*, classifying visual patterns directly by using raw images. However, limited by the scale of labeled training data and computational power, LeNet-5

as well as its variants [46] did not perform well on more complex vision tasks until recently.

To better train deep networks, Hinton *et al.* in 2006 made a breakthrough and introduced Deep Belief Networks (DBNs) to greedily train each layer of DBNs in an unsupervised manner [28]. And since then, researchers have developed more methods to overcome the difficulties in training CNN architectures. Particularly, AlexNet [41], as one of the milestones, was proposed by Krizhevsky *et al.* in 2012 and was successfully applied to large-scale image classification in the well-known ImageNet Challenge. AlexNet contains five convolutional layers followed by 3 fully-connected (*fc*) layers. Compared with LeNet-5, two novels components were introduced in AlexNet:

1. ReLUs (Rectified Linear Units) are utilized to replace the tanh units, which makes the training process several times faster;
2. Dropout is introduced and proved to be very effective in alleviating overfitting.

Inspired by AlexNet, several variants, including VGGNet [75], GoogLeNet [81] and ResNet [24], have been proposed to further improve the performance of CNNs on visual recognition tasks:

VGGNet [75] has two versions, *i.e.*, VGG16 and VGG19 models, which contain 16 and 19 layers respectively. VGGNet pushed the depth of CNN architecture from 8 layer as in AlexNet to 16-19 layers, which largely improves the discriminative power. In addition, by using very small (3×3) convolutional filters, VGGNet is capable of capturing details in the input images.

GoogLeNet [81] is inspired by the Hebbian principle with multi-scale processing and it contains 22 layers. A novel CNN architecture commonly referred to as Inception is proposed to increase both the depth and the width of CNN while maintaining an affordable computational cost. There are several extensions upon this work, including BN-Inception-V2 [82], Inception-V3 [82] and Inception-V4 [80].

ResNet [24], as one of the latest deep architectures, has remarkably increased the depth of CNN to 152 layers using deep residual layers with skip connections. ResNet won the 1st place in the 2015 ImageNet Challenge and has recently been extended to more than 1000 layers on the CIFAR-10 dataset [25].

From AlexNet, VGGNet, GoogLeNet to the more recent ResNet, one trend in the evolution of these architectures is to deepen the network. The increased depth facilities the network to better approximate the target function, generating better feature representations with higher discriminative power. In addition, various methods and strategies have been proposed from different aspects, including but not limited to, Maxout [19], DropConnect [95] and Batch Normalization [32], to ease the training of deep networks. Please refer to [22, 5] for a more detailed review.

2.2. Recurrent Neural Networks (RNNs)

The CNN architectures discussed above are all feed-forward neural networks whose connections do not form cycles, which is insufficient for sequence labeling. To better explore the temporal information of sequential data, recurrent connection structures have been introduced, leading to the emergence of RNN. Different from feed-forward neural networks, RNNs allow cyclical connections to form cycles, which thus enables a "memory" of previous inputs to persist in the network's internal state [20]. It has been pointed out that a finite-sized RNN with sigmoid activation functions can simulate a universal Turing machine [72].

The basic RNN block, at a time step t, accepts an external input vector $\boldsymbol{x}^{(t)} \in \mathbb{R}^n$ and generates an output vector $\boldsymbol{z}^{(t)} \in \mathbb{R}^m$ via a sequence of hidden states $\boldsymbol{h}^{(t)} \in \mathbb{R}^r$:

$$\begin{aligned} \boldsymbol{h}^{(t)} &= \sigma\left(W_x \boldsymbol{x}^{(t)} + W_h \boldsymbol{h}^{(t-1)} + \boldsymbol{b}_h\right) \\ \boldsymbol{z}^{(t)} &= \operatorname{softmax}\left(W_z \boldsymbol{h}^{(t)} + \boldsymbol{b}_z\right) \end{aligned} \tag{1}$$

where $W_x \in \mathbb{R}^{r \times n}$, $W_h \in \mathbb{R}^{r \times r}$, and $W_z \in \mathbb{R}^{m \times r}$ are weight matrices and $\boldsymbol{b}_h$ and $\boldsymbol{b}_z$ are biases. The σ is defined as sigmoid function $\sigma(x) = \frac{1}{1+e^{-x}}$ and $\operatorname{softmax}(\cdot)$ is the softmax function.

A problem with RNN is that it is not well capable of modeling long-range dependencies and is unable to store information about past inputs for a very long period [6], though one large enough RNN should, in principle, be able to approximate the sequences of arbitrary complexity. Specifically, two well-known issues – *vanishing* and *exploding* gradients exist in training RNNs: the vanishing gradient problem refers to the exponential shrinking of gradients magnitude as they are propagated back through time; and the exploding gradient problem refers to the explosion of long-term components due to the

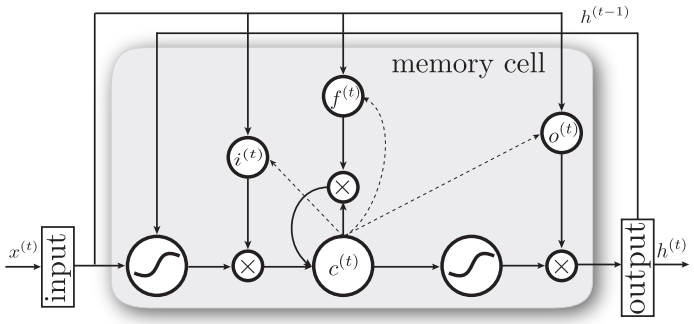

Figure 2: The structure of an LSTM unit. The figure is modified from [100].

large increase in the norm of the gradient during training sequences with long-term dependencies. To solve these issues, researchers introduced Long Short-Term Memory models.

Long Short-term Memory (LSTM) is an RNN variant, which was designed to store and access information in a long time sequence. Compared with standard RNNs, non-linear multiplicative gates and a memory cell are introduced. These gates, including input, output and forget gate, govern the information flow into and out of the memory cell. The structure of an LSTM unit is illustrated in Fig. 2.

More specifically, given a sequence of an external input vector $\boldsymbol{x}^{(t)} \in \mathbb{R}^n$, an LSTM maps the input to an output vector $\boldsymbol{z}^{(t)} \in \mathbb{R}^m$ by computing activations of the units in the network with the following equations recursively from $t = 1$ to $t = T$:

$$
\begin{aligned}
\boldsymbol{i}^{(t)} &= \sigma(W_{xi}\boldsymbol{x}^{(t)} + W_{hi}\boldsymbol{h}^{(t-1)} + W_{ci}\boldsymbol{c}^{(t)} + \boldsymbol{b}_i), \\
\boldsymbol{f}^{(t)} &= \sigma(W_{xf}\boldsymbol{x}^{(t)} + W_{hf}\boldsymbol{h}^{(t)} + W_{cf}\boldsymbol{c}^{(t)} + \boldsymbol{b}_f), \\
\boldsymbol{c}^{(t)} &= \boldsymbol{f}^{(t)}\boldsymbol{c}^{(t-1)} + \boldsymbol{i}_t \tanh(W_{xc}\boldsymbol{x}^{(t)} + W_{hc}\boldsymbol{h}^{(t-1)} + \boldsymbol{b}_c), \\
\boldsymbol{o}^{(t)} &= \sigma(W_{xo}\boldsymbol{x}^{(t)} + W_{ho}\boldsymbol{h}^{(t-1)} + W_{co}\boldsymbol{c}^{(t)} + \boldsymbol{b}_o), \\
\boldsymbol{h}^{(t)} &= \boldsymbol{o}^{(t)} \tanh(\boldsymbol{c}^{(t)}),
\end{aligned}
\tag{2}
$$

where $\boldsymbol{x}^{(t)}, \boldsymbol{h}^{(t)}$ are the input and hidden vectors with the subscription t denoting the t-th time step, $\boldsymbol{i}^{(t)}, \boldsymbol{f}^{(t)}, \boldsymbol{c}^{(t)}, \boldsymbol{o}^{(t)}$ are respectively the activation vectors of the input gate, forget gate, memory cell and output gate. $W_{\alpha\beta}$ denotes the weight matrix between α and β. For example, the weight matrix from the input $\boldsymbol{x}^{(t)}$ to the input gate $\boldsymbol{i}^{(t)}$ is W_{xi}.

In Eq (2) and Fig. (2) and at time step t, the input $\boldsymbol{x}^{(t)}$ and the previous states $\boldsymbol{h}^{(t-1)}$ are used as the input of LSTM. The information of memory cell is updated/controlled from two sources, *i.e.*, (1) the previous cell memory unit $\boldsymbol{c}^{(t-1)}$ and (2) the input gate's activation $\boldsymbol{i}_t$. Specifically, $\boldsymbol{c}^{(t-1)}$ is multiplied by the activation from the forget gate $\boldsymbol{f}^{(t)}$, which learns to forget the information of the previous states. In contrast, the $\boldsymbol{i}_t$ is combined with the new input signal to consider new information. Additionally, LSTM also utilizes the output gate $\boldsymbol{o}^{(t)}$ to control the information received by hidden state variable $\boldsymbol{h}^{(t)}$. To sum up, with these explicitly designed memory units and gates, LSTM is able to exploit the long-range temporal memory and avoids the issues of vanishing/exploding gradients. LSTM has recently been popularly used for video analysis, as will be discussed in the following sections.

3. Video Classification

The sheer volume of video data has motivated approaches to automatically categorizing video contents according to classes such as human activities and complex events. Extensive efforts have been devoted to design effective feature representations and robust classifiers. Recently, researchers have attempted to apply deep learning techniques to the video domain. In the following, we review two categories of deep learning algorithms for video classification, *i.e.*, supervised deep learning and unsupervised feature learning.

3.1. Supervised Deep Learning for Classification

3.1.1. Image-Based Video Classification

The great success of CNN features on image analysis tasks [18, 62] has stimulated the utilization of deep features for video classification. The general idea is to treat a video clip as a collection of frames, and then for each frame feature representation could be derived by running a feed-forward pass till a certain fully-connected layer with state-of-the-art deep models pre-trained on ImageNet [13], including AlexNet [41], VGGNet [75], GoogleNet [81] and ResNet [24], as discussed earlier. Finally, frame-level features are averaged into video-level representations as inputs of standard classifiers for recognition, such as the well-known SVMs.

Among the works on image-based video classification, Zha *et al.* systematically studied the performance of image-based video recognition using features from different layers of deep models together with multiple kernels

for classification [114]. They demonstrated that off-the-shelf CNN features coupled with kernel SVMs can obtain decent recognition performance. Motivated by the advanced feature encoding strategies in images [69], Xu *et al.* proposed to obtain video-level representation through VLAD encoding [33], which can attain performance gain over the trivial averaging pooling approach [105].

3.1.2. End-to-End CNN Architectures

The effectiveness of CNNs on a variety of tasks lies in its capability to learn features from raw data as an end-to-end pipeline targeting at a particular task [81, 50, 17]. Therefore, different from the image-based classification methods, there are many works focusing on applying CNN models to the video domain with an aim to learn hidden spatial-temporal patterns. Ji *et al.* introduced the 3D CNN model that operates on stacked video frames, extending the traditional 2D CNN designed for images to the spatial-temporal space [34]. The 3D CNN utilizes 3D kernels for convolution to learn motion information between adjacent frames in volumes segmented by human detectors. Karparthy *et al.* compared several similar architectures on a large scale video dataset in order to explore how to better extend the original CNN architectures to learn spatial-temporal clues in videos [39]. They found that the performance of the CNN model with a single frame as input achieves similar results to models operating on a stack of frames, and they also suggested that a mixed-resolution architecture consisting of a low-resolution context and a high-resolution stream could speed up training effectively. Recently, Tran *et al.* also utilized 3D convolutions with modern deep architectures [87]. However, they adopted full frames as the inputs of 3D CNNs instead of the segmented volumes in [34].

Though the extension of conventional CNN models by stacking frames makes sense, the performance of such models is worse than that of state-of-the-art hand-crafted features [96]. This may because the spatial-temporal patterns in videos are too complex to be captured by deep models with insufficient training data. In addition, the training of CNNs with inputs of 3D volumes is usually time-consuming. To effectively handle 3D signals, Sun *et al.* introduced factorized spatio-temporal convolutional networks that factorize the original 3D convolution kernel learning as a sequential process of learning 2D spatial kernels in the lower layer [78]. In addition, motived by the fact that videos can naturally be decomposed into spatial and temporal components, Simonyan and Zisserman proposed a two-stream approach (see

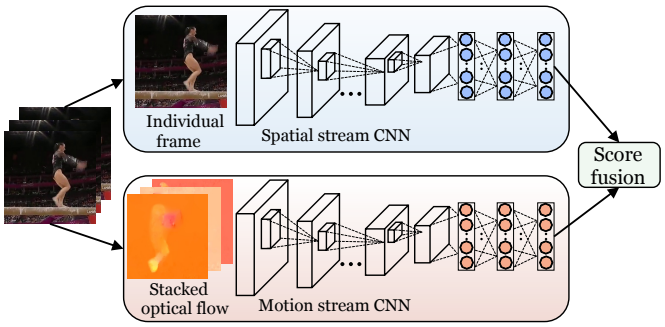

Figure 3: Two-stream CNN framework. Replicated from [101].

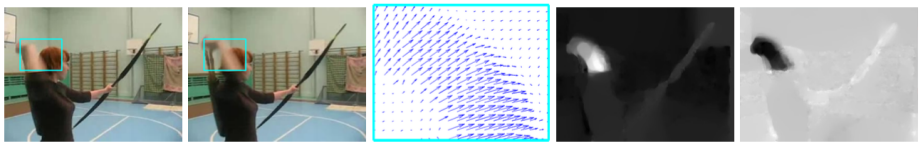

Figure 4: Examples of optical flow images. Replicated from [74].

Figure 3), which breaks down the learning of video representation into separate feature learning of spatial and temporal clues [74]. More specifically, the authors first adopted a typical spatial CNN to model appearance information with raw RGB frames as inputs. To account for temporal clues among adjacent frames, they explicitly generated multiple-frame dense optical flow, upon which a temporal CNN is trained. The dense optical flow is derived from computing displacement vector fields between adjacent frames (see Figure 4), which represent motions in an explicit way making the training of the network easier. Finally, at test time, each individual CNN generates a prediction by averaging scores from uniformly sampled 25 frames (optical flow frames) for a video clip, and then the final output is produced by the weighted sum of scores from the two streams. The authors reported promising results on two action recognition benchmarks. As the two-stream approach contains many implementation choices that may affect the performance, Ye *et al.* evaluated different options, including dropout ratio, network architecture, *etc.*, and discussed their findings in [112].

Very recently, there are several extensions of the two-stream approach. Wang *et al.* utilized the point trajectories from the improved dense tra-

jectories [96] to pool two-stream convolutional feature maps to generate Trajectory-Pooled Descriptors (TDD) [97]. Feichtenhofer *et al.* improved the two-stream approach by exploring a better fusion approach to combine spatial and temporal streams [15]. They found that two streams could be fused using convolutional layers rather than averaging classification scores to better model the correlations of spatial and temporal streams. Wang *et al.* introduced temporal segment networks, where each segment is used as the input of a two-stream network and the final prediction of a video clip is produced by a consensus function combining segment scores [98]. Zhang *et al.* proposed to replace the optical flow images with motion vectors with an aim to achieve real-time action recognition [115]. More recently, Wang *et al.* proposed to learn feature representation by modeling an action as a transformation between an initial state (condition) to a new state (effect) with two Siamese CNN networks, operating on RGB frames and optical flow images [99]. Similar to the original two-stream approach, they then fused the classification scores from two streams linearly to obtain final predictions. They reported better results on two challenging benchmarks than [74], possibly because the transformation from precondition to effect could implicitly model the temporal coherence in videos. Zhu *et al.* proposed a key volume mining approach that attempts to identify key volumes and perform classification at the same time [116]. Bilen *et al.* introduced the dynamic image to represent motions with rank pooling in videos, upon which a CNN model is trained for recognition [7].

3.1.3. Modeling Long-Term Temporal Dynamics

As discussed earlier, the temporal CNN in the two-stream approach [74] explicitly captures the motion information among adjacent frames, which however only depicts movements within a short time window. In addition, during the training of CNN models, each sweep takes a single frame (or a stacked optical frame image) as inputs of the network, failing to take the order of frames into account. This is not sufficient for video analysis, since complicated events/actions in videos usually consist of multiple actions happening over a long time. For instance, a "making pizza" event can be decomposed into several sequential actions, including "making the dough", "topping" and "baking". Therefore, researchers have recently attempted to leverage RNN models to account for the temporal dynamics in videos, among which LSTM is a good fit without suffering from the "vanishing gradient" effect and has demonstrated its effectiveness in several tasks like image/video

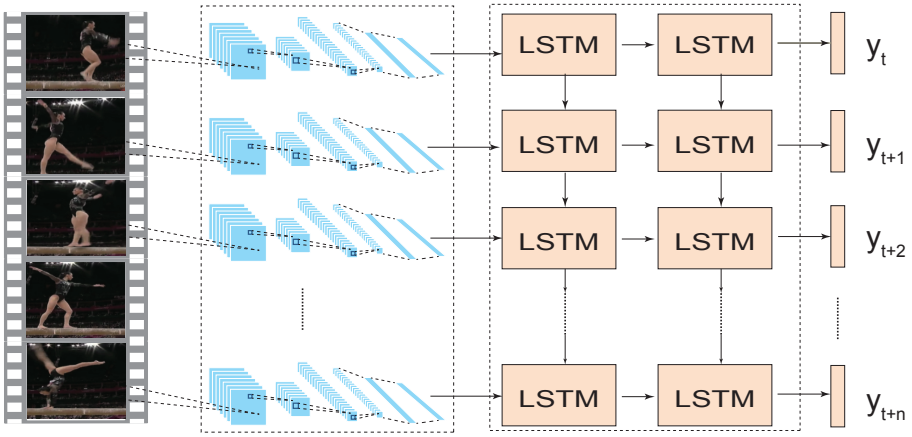

Figure 5: Utilizing LSTMs to explore temporal dynamics in videos with CNN features as inputs.

captioning [14, 106] (to be discussed in detail later) and speech analysis [21].

Donahue *et al.* trained two two-layer LSTM networks (Figure 5) for action recognition [14] with features from the two-stream approach. They also tried to fine-tune the CNN models together with LSTM but did not obtain significant performance gain compared with only training the LSTM model. Wu *et al.* fused the outputs of LSTM models with CNN models to jointly model spatial-temporal clues for video classification and observed that CNNs and LSTMs are highly complementary [101]. Ng *et al.* further trained a 5-layer LSTM model and compared several pooling strategies [54]. Interestingly, the deep LSTM model performs on par with single frame CNN on a large YouTube video dataset called Sports-1M, which may because the videos in this dataset are uploaded by ordinary users without professional editing and contain cluttered backgrounds and severe camera motion. Veeriah *et al.* introduced a differential gating scheme for LSTM to emphasize on the change in information gain [90] to remove redundancy in videos. Recently, in a multi-granular spatio-temporal architecture [47], LSTMs are utilized to further model the temporal information of frame, motion and clip streams. Besides, Wu *et al.* further employed a CNN operating on spectrograms derived from soundtracks of videos to compensate visual clues captured by CNN and LSTMs, and demonstrated strong results [100].

3.1.4. Incorporating Visual Attention

Videos contain many frames. Using all of them is computationally expensive and may degrade the performance of recognizing a class of interest as not all the frames are relevant. This issue has motivated researchers to leverage the attention mechanism to identify the most discriminative spatial-temporal volumes that are directly related to the targeting semantic class. Sharma *et al.* proposed the first attention LSTM for action recognition with a soft-attention mechanism to attach higher importance to the learned relevant parts in video frames [71]. More recently, Li *et al.* introduced the VideoLSTM, which applied attention in Convolutional LSTM models to discover relevant spatial-temporal volumes [48]. In addition to soft-attention, VideoLSTM also employed motion-based attention derived from optical flow images for better action localization.

3.2. Unsupervised Video Feature Learning

Current remarkable improvements with deep learning heavily rely on a large amount of labeled data. However, scaling up to thousands of video categories presents significant challenges due to insurmountable annotation efforts even at video level, not to mention frame level fine-grained labels. Therefore, the utilization of unsupervised learning, integrating spatial and temporal context information, is a promising way to find and represent structures in videos. Graham *et al.* proposed a Convolutional Gated Boltzmann Machine to learn to represent optical flow and describe motion [84]. Le *et al.* utilized two-layer ISA models to learn spatial-temporal models for action recognition [44]. More recently, Srivastava *et al.* adopted an encoder-decoder LSTM to learn feature representations in an unsupervised way [77]. They first mapped an input sequence into a fixed length representation by an encoder LSTM, which will be further decoded with single or multiple decoder LSTMs to perform different tasks, such as reconstructing the input sequence, or predicting the future sequence. The model is first pre-trained on YouTube data without manual labels, and then fine-tuned on standard benchmarks to recognize actions. Pan *et al.* explored both local temporal coherence and holistic graph structure preservation to learn a deep intrinsic video representation in an end-to-end fashion [57]. Nicolas *et al.* leveraged convolutional maps from different layers of a pre-trained CNN as the input of a GRU-RNN to learn video representations [3].

Summary: The latest developments discussed above have demonstrated the effectiveness of deep learning for video classification. However, current deep

learning approaches for video classification usually resort to popular deep models in image and speech domain. The complicated nature of video data, containing abundant spatial, temporal and acoustic clues, makes off-the-shelf deep models insufficient for video related tasks. This highlights the need for a tailored network to effectively capture spatial and acoustic information, and the most importantly to model temporal dynamics. In addition, training CNN/LSTM models requires manual labels which are usually expensive and time-consuming to acquire, and hence one promising direction is to make full utilization of the substantial amounts of unlabeled video data and rich contextual clues to derive better video representations.

4. Video Captioning

Video captioning is a new problem that has received increasing attention from both computer vision and natural language processing communities. Given an input video, the goal is to automatically generate a complete and natural sentence, ideally encapsulating the most informative dynamics in the video. In this section, we elaborate the problem by surveying the state-of-the-art methods. We classify existing methods in terms of different strategies for sentence modeling. In particular, we distill a common architecture of combining convolutional and recurrent neural networks for video captioning. As video captioning is an emerging area, we start from introducing the problem in detail.

4.1. Problem Introduction

Although there have already been extensive research on video tagging [73, 108] and image captioning [94, 14], video-level captioning has its own characteristics and thus is different from tagging and image/frame-level captioning. A video tag is usually the name of a specific object, action, or event, which is recognized in the video (*e.g.*, "baby," "boy," and "chair" in Figure 6). Image (frame) captioning goes beyond tagging by describing an image (frame) with a natural sentence, where the spatial relationships between objects or object and action are further described (*e.g.*, "Two baby boys are in the chair" generated on one single frame of Figure 6). Video captioning has been taken as an even more challenging problem, as the description generation model should be powerful enough not only to recognize scenes and objects in the video, but also be capable of modeling their spatio-temporal

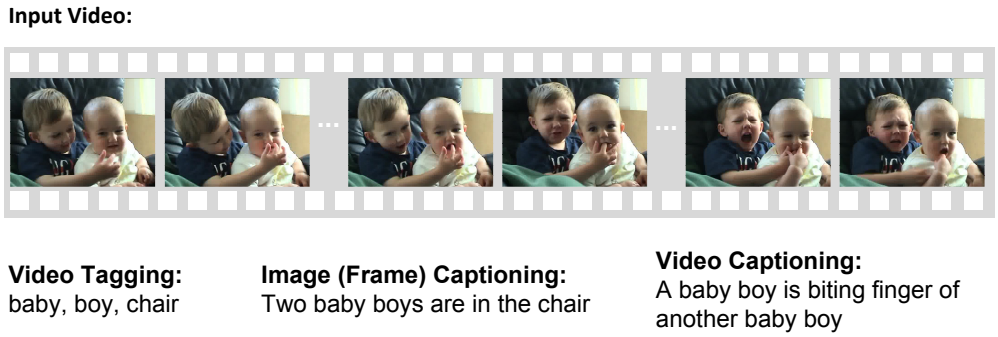
Input Video:
Video Tagging:
baby, boy, chair
Image (Frame) Captioning:
Two baby boys are in the chair
Video Captioning:
A baby boy is biting finger of another baby boy

Figure 6: Examples of video tagging, image (frame) captioning and video captioning. The input is a short video, while the output is a text response to this video, in the form of individual words (tags), a natural sentence describing one single image (frame) and dynamic video contents, respectively.

relationships and the dynamics expressed in a natural sentence (e.g., "A baby boy is biting finger of another baby boy" for the video in Figure 6).

Despite the difficulty of the problem, there have been several attempts to address video caption generation [58, 113, 103], which are mainly inspired by recent advances in machine translation [79]. The elegant recipes behind are the promising developments of the CNNs and the RNNs. In general, 2D [75] and/or 3D CNNs [87] are exploited to extract deep visual representations and LSTM [29] is utilized to generate the sentence word by word. More sophisticated frameworks by additionally integrating internal or external knowledge in the form of high-level semantic attributes or further exploring the relationship between the semantics of sentence and video content have also been studied for this problem.

In the following we start by presenting a comprehensive review of video captioning methods through two main categories based on the strategies for sentence generation in Sec. 4.2 and generalizing a common architecture by leveraging sequence learning for video captioning in Sec. 4.3.

4.2. Approaches for Video Captioning

There are mainly two directions for video captioning: template-based language model [40, 66, 63, 23, 104] and sequence learning models (*e.g.*, RNNs) [14, 58, 103, 113, 92, 106, 93, 93]. The former predefines the special rule for language grammar and splits sentence into several parts (e.g., subject, verb, object). With such sentence fragments, many works align each part with de-

tected words from visual content by object recognition and then generate a sentence with language constraints. The latter is to leverage sequence learning models to directly learn a translatable mapping between video content and sentence. We will review the state-of-the-art research along these two dimensions.

4.2.1. Template-based Language Model

Most of the approaches in this direction highly depend on the templates of sentence and always generate sentence with syntactical structure. [40] is one of the early works that built a concept hierarchy of actions for natural language description of human activities. Tan *et al.* [83] proposed to use predefined concepts and sentence templates for video event recounting. Rohrbach *et al.* learned a CRF to model the relationships between different components of the input video and generate descriptions for videos [66]. Furthermore, by incorporating semantic unaries and hand-centric features, Rohrbach *et al.* utilized CRF-based approach to generate coherent video descriptions [63]. In [23], Guadarrama *et al.* used semantic hierarchies to choose an appropriate level of the specificity and accuracy of sentence fragments. Recently, a deep joint video-language embedding model in [104] was designed for video sentence generation.

4.2.2. Sequence Learning

Different from template-based language model, sequence learning based methods can learn the probability distribution in the common space of visual content and textual sentence and generate novel sentence with more flexible syntactical structure. Donahua *et al.* employed a CRF to predict activity, object and location present in the video input. These representations are concatenated into an input sequence and then translated to a natural sentence with LSTM model [14]. Later in [93], Venugopalan *et al.* presented an end-to-end neural networks to generate video descriptions, which only reads the sequence of video frames. By mean pooling, the features over all the frames can be represented by one single vector, which are the input of the following LSTM model for sentence generation. The framework was then extended by inputting both frames and optical flow images into an encoder-decoder LSTM in [92]. Furthermore, inspired by the idea of learning visual-semantic embedding space in search [59, 107], Pan *et al.* additionally considered the relevance between sentence semantics and video content as a regularizer in LSTM based architecture [58]. Compared to mean pooling,

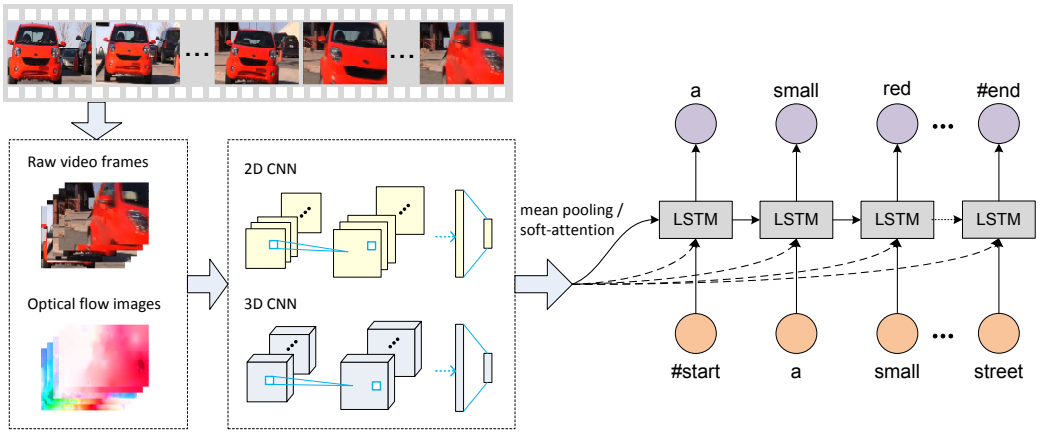

Figure 7: A common architecture to video captioning by sequence learning. The video representations are produced by mean pooling or soft-attention over the visual features of raw video frames/optical flow images/video clips, extracted by 2D/3D CNNs. The sentence is generated word by word in the following LSTM based on the video representations.

Yao *et al.* proposed to utilize the temporal attention mechanism to exploit temporal structure as well as a spatio-temporal convolutional neural network to obtain local action features. Then, the resulting video representations are fed into the text-generating RNN [106]. In addition, similar to the knowledge transfer from image domain to video domain in [110, 109], Liu *et al.* leveraged the learnt models on image captioning to generate a caption for each video frame and incorporate the obtained captions which are regarded as the attributes of each frame into a sequence to sequence architecture to generate video descriptions in [49].

4.3. A Common Architecture to Video Captioning

To better summarize the frameworks of video captioning by sequence learning, we illustrate a common architecture as shown in Figure 7. Given a video, 2D and/or 3D CNNs is utilized to extract visual features on raw video frames, optical flow images and video clips. The video-level representations are produced by mean pooling or soft attention over these visual features. Then, an LSTM is trained for generating video sentence based on the video-level representations.

Technically, suppose we have a video $\mathcal{V}$ with N_v sample frames/optical images/clips (uniform sampling) to be described by a textual sentence $\mathcal{S}$,

where $\mathcal{S} = \{w_1, w_2, ..., w_{N_s}\}$ consisting of N_s words. Let $\mathbf{v} \in \mathbb{R}^{D_v}$ and $\mathbf{w}_t \in \mathbb{R}^{D_w}$ denote the D_v-dimensional visual features of a video $\mathcal{V}$ and the D_w-dimensional textual features of the t-th word in sentence $\mathcal{S}$, respectively. As a sentence consists of a sequence of words, a sentence can be represented by a $D_w \times N_s$ matrix $\mathbf{W} \equiv [\mathbf{w}_1, \mathbf{w}_2, ..., \mathbf{w}_{N_s}]$, with each word in the sentence as its column vector. Hence, given the video representations $\mathbf{v}$, we aim to estimate the conditional probability of the output word sequence $\{w_1, w_2, ..., w_{N_s}\}$, *i.e.*,

$$\Pr(\mathbf{w}_1, \mathbf{w}_2, ..., \mathbf{w}_{N_s}|\mathbf{v}). \tag{3}$$

Since the model produces one word in the sentence at each time step, it is natural to apply chain rule to model the joint probability over the sequential words. Thus, the log probability of the sentence is given by the sum of the log probabilities over the words and can be expressed as:

$$\log \Pr(\mathbf{W}|\mathbf{v}) = \sum_{t=1}^{N_s} \log \Pr(\mathbf{w}_t|\, \mathbf{v}, \mathbf{w}_1, \ldots, \mathbf{w}_{t-1}). \tag{4}$$

In the model training, we feed the start sign word *#start* into LSTM, which indicates the starting of sentence generation process. We aim to maximize the log probability of the output video description $\mathcal{S}$ given the video representations, the pervious words it has seen and the model parameters θ, which can be formulated as

$$\theta^* = \underset{\theta}{argmax} \sum_{t=1}^{N_s} \log \Pr(\mathbf{w}_t|\, \mathbf{v}, \mathbf{w}_1, \ldots, \mathbf{w}_{t-1}; \theta). \tag{5}$$

This log probability is calculated and optimized over the whole training dataset using stochastic gradient descent. Note that the end sign word *#end* is required to terminate the description generation. When inference, we choose the word with maximum probability at each time step and set it as LSTM input for next time step until the end sign word is emitted.

Summary: The introduction of the video captioning problem is relatively new. Recently, this task sparks significant interests and may be regarded as the ultimate goal of video understanding. Video captioning is a complex problem and it is initially encouraged by the fundamental technological advances in recognition which could effectively recognize key objects or scenes from video contents. The developments of RNNs in machine translation further accelerated the growth of this research direction. The recent results,

although encouraging, are still indisputably far from practical use, as the forms of the generated sentences are simple and the vocabulary is still limited. How to generate free-form sentences and support open-vocabulary is vital to this task in the future.

5. Benchmarks and Competitions

We now discuss popular benchmarks and competitions for video classification in Sec. 5.1 and video captioning in Sec. 5.2.

5.1. Classification

Research on video classification has been stimulated largely by the release of the large and challenging video datasets such as UCF101 [76], HMDB51 [42] and FCVID [37] and the open competitions organized by fellow researchers, including the THUMOS challenge [36], the ActivityNet Large Scale Activity Recognition Challenge [26] and the TRECVID Multimedia Event Detection (MED) task [56]. In the following, we first discuss related datasets according to the list shown in Table 1 and then summarize the results of existing works.

5.1.1. Datasets

KTH dataset [70] is one of the earliest benchmarks for human action recognition. It contains 600 short videos of 6 human actions acted by 25 people in four different scenarios.

Weizmann datasets [8] is another very early and simple dataset, consisting of 81 short videos associated with 9 actions performed by 9 actors.

Kodak Consumer Videos dataset [51] was recorded by around 100 customers of the Eastman Kodak Company. The dataset collected $1,358$ video clips labeled with 25 concepts (including activities, scenes and single objects) as a part of the Kodak concept ontology.

Hollywood Human Action dataset [43] contains 8 action classes collected from 32 Hollywood movies totaling 430 video clips. It was further extended to the **Hollywood2** [52] dataset, which is composed of 12 actions of 69 Hollywood movies with $1,707$ video clips in total. The Hollywood series are challenging due to cluttered background and severe camera motion that widely exist in this dataset.

Dataset	#Video	#Class	Released Year	Background
KTH	600	6	2004	Clean Static
Weizmann	81	9	2005	Clean Static
Kodak	1,358	25	2007	Dynamic
Hollywood	430	8	2008	Dynamic
Hollywood2	1,787	12	2009	Dynamic
MCG-WEBV	234,414	15	2009	Dynamic
Olympic Sports	800	16	2010	Dynamic
HMDB51	6,766	51	2011	Dynamic
CCV	9,317	20	2011	Dynamic
UCF-101	13,320	101	2012	Dynamic
THUMOS-2014	18,394	101	2014	Dynamic
MED-2014 (Dev. set)	≈31,000	20	2014	Dynamic
Sports-1M	1,133,158	487	2014	Dynamic
ActivityNet	27,901	203	2015	Dynamic
EventNet	95,321	500	2015	Dynamic
MPII Human Pose	20,943	410	2014	Dynamic
FCVID	91,223	239	2015	Dynamic

Table 1: Popular benchmark datasets for video classification, sorted by the year of construction.

MCG-WEBV dataset [9] is another large set of YouTube videos which has 234,414 web videos with annotations on several topic-level events like "a conflict at Gaza".

Olympic Sports [55] includes 800 video clips and 16 action classes. It was first introduced in 2010 and, different from the previous datasets, all the videos were downloaded from the Internet.

HMDB51 dataset [42] is comprised of 6,766 videos annotated into 51 classes. The videos are from a variety of sources, including movies and YouTube consumer videos.

Columbia Consumer Videos (CCV) dataset [38] was constructed in 2011, aiming to stimulate the research on Internet consumer video analysis. It contains 9,317 user generated videos from YouTube, which were annotated into 20 classes, including objects (*e.g.*, "cat" and "dog"), scenes (*e.g.*, "beach" and "playground"), sports events (*e.g.*, "basket-

ball" and "soccer") and social activities (*e.g.*, "birthday" and "graduation").

UCF-101 & THUMOS-2014 dataset [76] is another popular benchmark for human action recognition in videos, which consists of 13,320 video clips (27 hours in total) with 101 annotated classes. More recently, the THUMOS-2014 Action Recognition Challenge [36] created a benchmark by extending upon the UCF-101 dataset (used as the training set). Additional videos were collected from the Internet, including 2,500 background videos, 1,000 validation videos and 1,574 test videos.

TRECVID MED dataset [56] was released and annually updated by the task of MED created by NIST since 2010. Each year an extended dataset based on datasets from competitions of previous years is constructed and released for worldwide system comparison. For example, in 2014 the MED dataset contains 20 events, such as "birthday party" ,"bike trick" , *etc.* According to NIST, in the development set, there are around 8,000 videos for training and 23,000 videos as dry-run validation samples (1,200 hours in total). The MED dataset is only available to the participants of the task, and the labels of the official test set (200,000 videos) is not available even to the participants.

Sports-1M dataset [39] consists of 1 million YouTube videos in 487 classes, such as "bowling", "cycling", "rafting", and *etc.* and is available in 2014. The video annotations were automatically derived by analyzing online textual contexts of the videos. Therefore the labels of this dataset is not clean, but the authors claimed that the quality of annotation is fairly good.

ActivityNet dataset [26] is another large-scale video dataset for human activity recognition and understanding and was released in 2015. It consists of 27,801 video clips annotated into 203 activity classes, totaling 849 hours of video. Compared with existing datasets, ActivityNet contains more fine-grained action categories (*e.g.*, "drinking beer" and "drinking coffee").

EventNet dataset [111] consists of 500 events and 4,490 event-specific concepts and was released in 2015. It includes automatic detection models for the video events and some constituent concepts with around

95, 000 training videos from YouTube. Similar to Sports-1M, EventNet was labeled by online textual information rather than manually labeled.

MPII Human Pose dataset [2] includes around $25K$ images containing over $40K$ people with annotated body joints. According to an established taxonomy of human activities (410 in total), the collected images (from YouTube videos) were provided with activity labels.

Fudan-Columbia Video Dataset (FCVID) dataset [37] contains 91, 223 web videos annotated manually into 239 categories. The categories cover a wide range of topics, such as social events (*e.g.*, "tailgate party"), procedural events (*e.g.*, "making cake"), object appearances (*e.g.*, "panda") and scenes (*e.g.*, "beach").

5.1.2. Competitions

To advance the state of the arts in video classification, several competitions have been introduced with the aim of exploring and evaluating new approaches in realistic settings. We briefly introduce three representative challenges here.

THUMOS challenge [31] was firstly introduced in 2013 in the computer vision community, aiming to explore and evaluate new approaches for large-scale action recognition of Internet videos. The three editions of the challenge organized in 2013–2015 had made THUMOS a common benchmark for action classification and detection. See [31] for details.

TRECVID Multimedia Event Detection (MED) Task [56] aims at detecting whether a video clip contains an instance of a specific event. Specifically, based on the released TRECVID MED dataset each year, each participant is required to provide for each testing video the confidence score of how likely one particular event happens in this video. 20 pre-specified events are used each year; and this task adopts the metrics of Average Precision (AP) and inferred AP for event detection. Each event was also complemented with an event kit, *i.e.* the textual description of the event as well as the potentially useful information about related concepts that are likely contained in the event.

The ActivityNet Large Scale Activity Recognition Challenge [26] was first organized as a workshop in 2016. This challenge is based on the

Methods	UCF-101	HMDB51
LRCN [14]	82.9	-
LSTM-composite [77]	84.3	-
$F_{ST}CN$ [78]	88.1	59.1
C3D [87]	86.7	-
Two-Stream [74]	88.0	59.4
LSTM [54]	88.6	-
Image-Based [114]	89.6	-
Transformation CNN [99]	92.4	63.4
Multi-Stream [100]	92.6	-
Key Volume Mining [116]	92.7	67.2
Convolutional Two-Stream [15]	93.5	69.2
Temporal Segment Networks [98]	94.2	69.4

Table 2: Comparison of recent video classification methods on UCF-101 and HMDB51 datasets.

ActivityNet dataset [26]; with the aim of recognizing high-level and goal oriented activities. By using 203 activity categories, there are two tasks in this challenge: (1) Untrimmed Classification Challenge; (2) Detection Challenge, which is to predict the labels and temporal extents of the activities present in videos.

5.1.3. Results of Existing Methods

Some of the datasets introduced above have been popularly adopted in the literature. We summarize the results of several recent approaches on UCF-101 and HMDB51 in Table 2, where we can see the fast pace of development in this area. Results on video classification are mostly measured by the AP (for a single class) and mean AP (for multiple classes), which are not introduced in detail as they are well-known.

5.2. Captioning

A number of datasets have been proposed for video captioning, which commonly contain pairs of a video and its corresponding sentences annotated by humans. This section summarizes the existing datasets and the adopted evaluation metrics, followed by quantitative results of representative methods.

Dataset	Context	Sentence Source	#Video	#Clip	#Sentence	#Word
MSVD	multi-category	AMT workers	–	1,970	70,028	607,339
YouCook	cooking	AMT workers	88	–	2,668	42,457
TACoS-ML	cooking	AMT workers	273	14,105	52,593	–
M-VAD	movie	DVS	92	48,986	55,905	519,933
MPII-MD	movie	Script+DVS	94	68,337	68,375	653,467
MSR-VTT-10K	20 categories	AMT workers	7,180	10,000	200,000	1,856,523

Table 3: Recent benchmark datasets for video captioning.

5.2.1. Datasets

Table 3 summarizes key statistics and comparisons of popular datasets for video captioning. Figure 8 shows a few examples from some of the datasets.

Microsoft Research Video Description Corpus (MSVD) [10] contains 1,970 YouTube snippets collected on Amazon Mechanical Turk (AMT) by requesting workers to pick short clips depicting a single activity. The video clips are then labeled with single sentence descriptions by annotators. The original corpus has multi-lingual descriptions and only the English descriptions are commonly exploited on video captioning task. Specifically, there are roughly 40 available English descriptions per video and the standard split of MSVD is 1,200 videos for training, 100 for validation and 670 for testing, as suggested in [23].

YouCook dataset [12] consists of 88 in-house cooking videos crawled from YouTube and is roughly uniformly split into six different cooking styles, such as baking and grilling. All the videos are taken in a third-view viewpoint and different kitchen environments. Each video is annotated with multiple human descriptions by AMT. Each annotator in AMT is instructed to describe the video in at least three sentences totaling a minimum of 15 words, resulting in 2,668 sentences for all the videos.

TACoS Multi-Level Corpus (TACoS-ML) [63] is mainly built based on MPII Cooking Activities dataset 2.0 [67] which records different activities in the cooking scenario. TACoS-ML consists of 185 long videos and the text descriptions are collected via AMT workers. Each AMT worker annotates a sequence of temporal intervals across the long video, pairing every interval with a single short sentence. There are 14,105 distinct intervals and 52,593 sentences in total.

(a) MSVD dataset

Sentences:

- a dog walks around on its front legs
- the dog is doing a handstand
- a pug is trying for balance walk on its two legs

Sentences:

- a man lights a match book on fire
- a man playing with fire sticks
- a man lights matches and yells

(b) M-VAD dataset

Sentence:

- Later he drags someone through a jog.

Sentence:

- A waiter brings a pastry with a candle.

(c) MPII-MD dataset

Sentence:

- He places his hands around her waist as she open her eyes.

Sentence:

- Someone' car is stopped by a couple of uniformed police.

(d) MSR-VTT-10K dataset

Sentences:

- people practising volley ball game in the play ground
- a man hitting a ball and he fallen
- a man is playing a football game on green land

Sentences:

- a cat is hanging out in a bassinet with a baby
- the cat is in the baby bed with the baby
- a cat plays with a child in a crib

Figure 8: Examples from (a) MSVD, (b) M-VAD, (c) MPII-MD and (d) MSR-VTT-10K datasets.

Montreal Video Annotation Dataset (M-VAD) [86] is composed of about 49,000 DVD movie snippets, which are extracted from 92 DVD movies. Each movie clip is accompanied with one single sentence from semi-automatically transcribed descriptive video service (DVS) narra-

Table 4: Reported results on the MSVD dataset, where B@N, M, and C are short for BLEU@N, METEOR, and CIDEr-D scores, respectively. All values are reported as percentage (%).

Methods	B@1	B@2	B@3	B@4	M	C
FGM [85]	-	-	-	13.7	23.9	-
LSTM-YT [93]	-	-	-	33.3	29.1	-
MM-VDN [102]	-	-	-	37.6	29.0	-
S2VT [92]	-	-	-	-	29.8	-
S2FT [49]	-	-	-	-	29.9	-
SA [106]	80.0	64.7	52.6	41.9	29.6	51.7
Glove+Deep Fusion [91]	-	-	-	42.1	31.4	-
LSTM-E [58]	78.8	66.0	55.4	45.3	31.0	-
GRU-RCN [3]	-	-	-	43.3	31.6	68.0
h-RNN [113]	81.5	70.4	60.4	49.9	32.6	65.8

tions. The fact that movies are always with high diversity of visual and textual content, and there is only one single reference sentence for each movie clip has made the video captioning task on M-VAD dataset very challenging.

MPII Movie Description Corpus (MPII-MD) [65] is another recent collection of movie description dataset, which is similar to M-VAD. It contains around 68,000 movie snippets from 94 Hollywood movies and each snippet is equipped with one single sentence from movie scripts and DVS.

MSR Video to Text (MSR-VTT-10K) [103] is the most recent large-scale benchmark for video captioning task, which contains 10K Web video clips of 41.2 hours, covering the most comprehensive 20 categories obtained from a commercial video search engine, e.g., music, people, gaming, sports, and TV shows. Each clip is annotated with about 20 natural sentences by AMT workers. The training/validation/test split is provided by the authors with 6,513 clips for training, 2,990 for validation and 497 for testing.

5.2.2. Evaluation Metrics

For quantitative evaluation of the video captioning task, three metrics are commonly adopted: BLEU@N [60], METEOR [4], and CIDEr [89]. Specifically, BLEU@N is a popular machine translation metric which measures the

Table 5: Reported results on (a) M-VAD and (b) MPII-MD datasets, where M is short for METEOR. All values are reported as percentage (%).

(a) M-VAD dataset.

Methods	M
SA [106]	4.3
Mean Pool [92]	6.1
Visual-Labels [64]	6.4
S2VT [92]	6.7
Glove+Deep Fusion [91]	6.8
LSTM-E [58]	6.7

(b) MPII-MD dataset.

Methods	M
SMT [65]	5.6
Mean Pool [92]	6.7
Visual-Labels [64]	7.0
S2VT [92]	7.1
Glove+Deep Fusion [91]	6.8
LSTM-E [58]	7.3

Table 6: Reported results on the TACoS-ML dataset, where B@N, M, and C are short for BLEU@N, METEOR, and CIDEr-D scores, respectively. All values are reported as percentage (%).

Methods	B@1	B@2	B@3	B@4	M	C
CRF-T [66]	56.4	44.7	33.2	25.3	26.0	124.8
CRF-M [63]	58.4	46.7	35.2	27.3	27.2	134.7
LRCN [14]	59.3	48.2	37.0	29.2	28.2	153.4
h-RNN [113]	60.8	49.6	38.5	30.5	28.7	160.2

fraction of N-gram (up to 4-gram) that are in common between a hypothesis and a reference or set of references. However, as pointed in [11], the N-gram matches for a high N (e.g., 4) rarely occur at a sentence-level, resulting in poor performance of BLEU@N especially when comparing individual sentences. Hence, another more effective evaluation metric METEOR is utilized along with BLEU@N, which is also widely used in NLP community. Different from BLEU@N, METEOR computes unigram precision and recall,

extending exact word matches to include similar words based on WordNet synonyms and stemmed tokens. Another important metric for image/video captioning is CIDEr, which measures consensus in image/video captioning by performing a Term Frequency Inverse Document Frequency (TF-IDF) weighting for each N-gram.

5.2.3. Results of Existing Methods

Most popular methods of video captioning have been evaluated on MSVD [10], M-VAD [86], MPII-MD [65], and TACoS-ML [63] datasets. We summarize the results on these four datasets in Table 4, 5, and 6. As can be seen, most of the works are very recent, indicating that video captioning is an emerging and fast-developing research topic.

6. Conclusion

In this paper, we have reviewed state-of-the-art deep learning techniques on two key topics related to video analysis, video classification and video captioning, both of which rely on the modeling of the abundant spatial and temporal information in videos. The essence of video classification with deep learning is to derive robust and discriminative feature representations from raw data through exploiting the substantial amounts of videos with an aim to achieve effective and efficient recognition, which could hence serve as a fundamental component in video captioning. Video captioning, on the other hand, focuses on bridging visual understanding and language description by joint modeling. We also provided a review of popular benchmarks and competitions for both video classification and captioning tasks. As deep learning for video analysis is an emerging and vibrant research area, we hope this paper could provide helpful information for both current and new researchers.

References

[1] Aggarwal, J. K., Ryoo, M. S., 2011. Human activity analysis: A review. ACM Computing Surveys.

[2] Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B., 2014. 2d human pose estimation: New benchmark and state of the art analysis. In: CVPR.

[3] Ballas, N., Yao, L., Pal, C., Courville, A., 2016. Delving deeper into convolutional networks for learning video representations. In: ICLR.

[4] Banerjee, S., Lavie, A., 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In: ACL Workshop.

[5] Bengio, Y., Courville, A., Vincent, P., 2013. Representation learning: A review and new perspectives. IEEE TPAMI.

[6] Bengio, Y., Simard, P., Frasconi, P., 1994. Learning long-term dependencies with gradient descent is difficult. IEEE TNN.

[7] Bilen, H., Fernando, B., Gavves, E., Vedaldi, A., Gould, S., 2016. Dynamic image networks for action recognition. In: CVPR.

[8] Blank, M., Gorelick, L., Shechtman, E., Irani, M., Basri, R., 2005. Actions as space-time shapes. In: ICCV.

[9] Cao, J., Zhang, Y.-D., Song, Y.-C., Chen, Z.-N., Zhang, X., Li, J.-T., 2009. MCG-WEBV: A benchmark dataset for web video analysis. Technical Report, CAS Institute of Computing Technology.

[10] Chen, D. L., Dolan, W. B., 2011. Collecting highly parallel data for paraphrase evaluation. In: ACL.

[11] Chen, X., Fang, H., Lin, T.-Y., Vedantam, R., Gupta, S., Dollár, P., Zitnick, C. L., 2015. Microsoft COCO captions: Data collection and evaluation server. arXiv preprint arXiv:1504.00325.

[12] Das, P., Xu, C., Doell, R. F., Corso, J. J., 2013. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In: CVPR.

[13] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. ImageNet: A large-scale hierarchical image database. In: CVPR.

[14] Donahue, J., Hendricks, L. A., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., Darrell, T., 2015. Long-term recurrent convolutional networks for visual recognition and description. In: CVPR.

[15] Feichtenhofer, C., Pinz, A., Zisserman, A., 2016. Convolutional two-stream network fusion for video action recognition. In: CVPR.

[16] Fukushima, K., 1980. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. Biological cybernetics.

[17] Girshick, R., 2015. Fast r-cnn. In: ICCV.

[18] Girshick, R., Donahue, J., Darrell, T., Malik, J., 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In: CVPR.

[19] Goodfellow, I. J., Warde-Farley, D., Mirza, M., Courville, A. C., Bengio, Y., 2013. Maxout networks. In: ICML.

[20] Graves, A., 2012. Supervised sequence labelling with recurrent neural networks. Springer.

[21] Graves, A., Mohamed, A., Hinton, G. E., 2013. Speech recognition with deep recurrent neural networks. In: ICASSP.

[22] Gu, J., Wang, Z., Kuen, J., Ma, L., Shahroudy, A., Shuai, B., Liu, T., Wang, X., Wang, G., 2016. Recent advances in convolutional neural networks. In: arXiv preprint arXiv:1512.07108.

[23] Guadarrama, S., Krishnamoorthy, N., Malkarnenkar, G., Venugopalan, S., Mooney, R., Darrell, T., Saenko, K., 2013. Youtube2text: Recognizing and describing arbitrary activities using semantic hierarchies and zero-shot recognition. In: ICCV.

[24] He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: CVPR.

[25] He, K., Zhang, X., Ren, S., Sun, J., 2016. Identity mappings in deep residual networks. In: ECCV.

[26] Heilbron, F. C., Escorcia, V., Ghanem, B., Niebles, J. C., 2015. Activitynet: A large-scale video benchmark for human activity understanding. In: CVPR.

[27] Hinton, G., Deng, L., Yu, D., Dahl, G. E., Mohamed, A.-r., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T. N., et al., 2012. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. IEEE Signal Processing Magazine.

[28] Hinton, G. E., Osindero, S., Teh, Y.-W., 2006. A fast learning algorithm for deep belief nets. Neural computation.

[29] Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. Neural computation.

[30] Hubel, D. H., Wiesel, T. N., 1968. Receptive fields and functional architecture of monkey striate cortex. The Journal of physiology.

[31] Idrees, H., Zamir, A. R., Jiang, Y.-G., Gorban, A., Laptev, I., Sukthankar, R., Shah, M., 2016. The thumos challenge on action recognition for videos "in the wild". arXiv preprint arXiv:1604.06182.

[32] Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML.

[33] Jégou, H., Douze, M., Schmid, C., Pérez, P., 2010. Aggregating local descriptors into a compact image representation. In: CVPR.

[34] Ji, S., Xu, W., Yang, M., Yu, K., 2010. 3d convolutional neural networks for human action recognition. In: ICML.

[35] Jiang, Y.-G., Bhattacharya, S., Chang, S.-F., Shah, M., 2013. High-level event recognition in unconstrained videos. IJMIR.

[36] Jiang, Y.-G., Liu, J., Roshan Zamir, A., Toderici, G., Laptev, I., Shah, M., Sukthankar, R., 2014. THUMOS challenge: Action recognition with a large number of classes. http://crcv.ucf.edu/THUMOS14/.

[37] Jiang, Y.-G., Wu, Z., Wang, J., Xue, X., Chang, S.-F., 2015. Exploiting feature and class relationships in video categorization with regularized deep neural networks. arXiv preprint arXiv:1502.07209.

[38] Jiang, Y.-G., Ye, G., Chang, S.-F., Ellis, D., Loui, A. C., 2011. Consumer video understanding: A benchmark database and an evaluation of human and machine performance. In: ACM ICMR.

[39] Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L., 2014. Large-scale video classification with convolutional neural networks. In: CVPR.

[40] Kojima, A., Tamura, T., Fukunaga, K., 2002. Natural language description of human activities from video images based on concept hierarchy of actions. IJCV.

[41] Krizhevsky, A., Sutskever, I., Hinton, G. E., 2012. Imagenet classification with deep convolutional neural networks. In: NIPS.

[42] Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T., 2011. Hmdb: a large video database for human motion recognition. In: ICCV.

[43] Laptev, I., Marszalek, M., Schmid, C., Rozenfeld, B., 2008. Learning realistic human actions from movies. In: CVPR.

[44] Le, Q. V., Zou, W. Y., Yeung, S. Y., Ng, A. Y., 2011. Learning hierarchical invariant spatio-temporal features for action recognition with independent subspace analysis. In: CVPR.

[45] LeCun, B. B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., Jackel, L. D., 1990. Handwritten digit recognition with a back-propagation network. In: NIPS.

[46] LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., 2001. Gradient-based learning applied to document recognition. In: Intelligent Signal Processing.

[47] Li, Q., Qiu, Z., Yao, T., Mei, T., Rui, Y., Luo, J., 2016. Action recognition by learning deep multi-granular spatio-temporal video representation. In: ICMR.

[48] Li, Z., Gavves, E., Jain, M., Snoek, C. G., 2016. Videolstm convolves, attends and flows for action recognition. arXiv preprint arXiv:1607.01794.

[49] Liu, Y., Shi, Z., 2016. Boosting video description generation by explicitly translating from frame-level captions. In: ACM Multimedia.

[50] Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation. In: CVPR.

[51] Loui, A., Luo, J., Chang, S.-F., Ellis, D., Jiang, W., Kennedy, L., Lee, K., Yanagawa, A., 2007. Kodak's consumer video benchmark data set: concept definition and annotation. In: ACM MIR Workshop.

[52] Marszalek, M., Laptev, I., Schmid, C., 2009. Actions in context. In: CVPR.

[53] McCulloch, W., Pitts, W., 1943. A logical calculus of the ideas immanent in nervous activity. Bulletin of Mathematical Biophysics.

[54] Ng, J. Y.-H., Hausknecht, M., Vijayanarasimhan, S., Vinyals, O., Monga, R., Toderici, G., 2015. Beyond short snippets: Deep networks for video classification. In: CVPR.

[55] Niebles, J. C., Chen, C.-W., Fei-Fei, L., 2010. Modeling temporal structure of decomposable motion segments for activity classification. In: ECCV.

[56] Over, P., Awad, G., Michel, M., Fiscus, J., Sanders, G., Kraaij, W., Smeaton, A. F., Quéenot, G., 2014. Trecvid 2014 – an overview of the goals, tasks, data, evaluation mechanisms and metrics. In: TRECVID Workshop.

[57] Pan, Y., Li, Y., Yao, T., Mei, T., Li, H., Rui, Y., 2016. Learning deep intrinsic video representation by exploring temporal coherence and graph structure. In: IJCAI.

[58] Pan, Y., Mei, T., Yao, T., Li, H., Rui, Y., 2016. Jointly modeling embedding and translation to bridge video and language. In: CVPR.

[59] Pan, Y., Yao, T., Mei, T., Li, H., Ngo, C.-W., Rui, Y., 2014. Click-through-based cross-view learning for image search. In: SIGIR.

[60] Papineni, K., Roukos, S., Ward, T., Zhu, W.-J., 2002. Bleu: a method for automatic evaluation of machine translation. In: ACL.

[61] Poppe, R., 2010. A survey on vision-based human action recognition. Image and vision computing.

[62] Razavian, A. S., Azizpour, H., Sullivan, J., Carlsson, S., 2014. CNN features off-the-shelf: an astounding baseline for recognition. In: CVPR Workshop.

[63] Rohrbach, A., Rohrbach, M., Qiu, W., Friedrich, A., Pinkal, M., Schiele, B., 2014. Coherent multi-sentence video description with variable level of detail. In: GCPR.

[64] Rohrbach, A., Rohrbach, M., Schiele, B., 2015. The long-short story of movie description. In: GCPR.

[65] Rohrbach, A., Rohrbach, M., Tandon, N., Schiele, B., 2015. A dataset for movie description. In: CVPR.

[66] Rohrbach, M., Qiu, W., Titov, I., Thater, S., Pinkal, M., Schiele, B., 2013. Translating video content to natural language descriptions. In: ICCV.

[67] Rohrbach, M., Rohrbach, A., Regneri, M., Amin, S., Andriluka, M., Pinkal, M., Schiele, B., 2015. Recognizing fine-grained and composite activities using hand-centric features and script data. IJCV.

[68] Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al., 2015. Imagenet large scale visual recognition challenge. IJCV.

[69] Sánchez, J., Perronnin, F., Mensink, T., Verbeek, J., 2013. Image classification with the fisher vector: Theory and practice. IJCV.

[70] Schuldt, C., Laptev, I., Caputo, B., 2004. Recognizing human actions: a local svm approach. In: ICPR.

[71] Sharma, S., Kiros, R., Salakhutdinov, R., 2015. Action recognition using visual attention. arXiv preprint arXiv:1511.04119.

[72] Siegelmann, H. T., Sontag, E. D., 1991. Turing computability with neural nets. Applied Mathematics Letters 6.

[73] Siersdorfer, S., Pedro, J. S., Sanderson, M., 2009. Automatic video tagging using content redundancy. In: SIGIR.

[74] Simonyan, K., Zisserman, A., 2014. Two-stream convolutional networks for action recognition in videos. In: NIPS.

[75] Simonyan, K., Zisserman, A., 2015. Very deep convolutional networks for large-scale image recognition. In: ICLR.

[76] Soomro, K., Zamir, A. R., Shah, M., 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. CoRR.

[77] Srivastava, N., Mansimov, E., Salakhutdinov, R., 2015. Unsupervised learning of video representations using LSTMs. In: ICML.

[78] Sun, L., Jia, K., Yeung, D.-Y., Shi, B. E., 2015. Human action recognition using factorized spatio-temporal convolutional networks. In: CVPR.

[79] Sutskever, I., Vinyals, O., Le, Q. V., 2014. Sequence to sequence learning with neural networks. In: NIPS.

[80] Szegedy, C., Ioffe, S., Vanhoucke, V., 2016. Inception-v4, inception-resnet and the impact of residual connections on learning. In: ICLR.

[81] Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions. In: CVPR.

[82] Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2015. Rethinking the inception architecture for computer vision. arXiv preprint arXiv:1512.00567.

[83] Tan, C.-C., Jiang, Y.-G., Ngo, C.-W., 2011. Towards textually describing complex video contents with audio-visual concept classifiers. In: ACM Multimedia.

[84] Taylor, G. W., Fergus, R., LeCun, Y., Bregler, C., 2010. Convolutional learning of spatio-temporal features. In: ECCV.

[85] Thomason, J., Venugopalan, S., Guadarrama, S., Saenko, K., Mooney, R., 2014. Integrating language and vision to generate natural language descriptions of videos in the wild. In: COLING.

[86] Torabi, A., Pal, C., Larochelle, H., Courville, A., 2015. Using descriptive video services to create a large data source for video annotation research. arXiv preprint arXiv:1503.01070.

[87] Tran, D., Bourdev, L. D., Fergus, R., Torresani, L., Paluri, M., 2015. C3d: Generic features for video analysis. In: ICCV.

[88] Turaga, P., Chellappa, R., Subrahmanian, V. S., Udrea, O., 2008. Machine recognition of human activities: A survey. IEEE TPAMI.

[89] Vedantam, R., Lawrence Zitnick, C., Parikh, D., 2015. Cider: Consensus-based image description evaluation. In: CVPR.

[90] Veeriah, V., Zhuang, N., Qi, G.-J., 2015. Differential recurrent neural networks for action recognition. In: ICCV.

[91] Venugopalan, S., Hendricks, L. A., Mooney, R., Saenko, K., 2016. Improving lstm-based video description with linguistic knowledge mined from text. arXiv preprint arXiv:1604.01729.

[92] Venugopalan, S., Rohrbach, M., Donahue, J., Mooney, R., Darrell, T., Saenko, K., 2015. Sequence to sequence - video to text. In: ICCV.

[93] Venugopalan, S., Xu, H., Donahue, J., Rohrbach, M., Mooney, R., Saenko, K., 2015. Translating videos to natural language using deep recurrent neural networks. In: NAACL HLT.

[94] Vinyals, O., Toshev, A., Bengio, S., Erhan, D., 2015. Show and tell: A neural image caption generator. In: CVPR.

[95] Wan, L., Zeiler, M., Zhang, S., Cun, Y. L., Fergus, R., 2013. Regularization of neural networks using dropconnect. In: ICML.

[96] Wang, H., Schmid, C., 2013. Action recognition with improved trajectories. In: ICCV.

[97] Wang, L., Qiao, Y., Tang, X., 2015. Action recognition with trajectory-pooled deep-convolutional descriptors. In: CVPR.

[98] Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Van Gool, L., 2016. Temporal segment networks: Towards good practices for deep action recognition. In: ECCV.

[99] Wang, X., Farhadi, A., Gupta, A., 2016. Actions ˜ transformations. In: CVPR.

[100] Wu, Z., Jiang, Y.-G., Wang, X., Ye, H., Xue, X., 2016. Multi-stream multi-class fusion of deep networks for video classification. In: ACM Multimedia.

[101] Wu, Z., Wang, X., Jiang, Y.-G., Ye, H., Xue, X., 2015. Modeling spatial-temporal clues in a hybrid deep learning framework for video classification. In: ACM Multimedia.

[102] Xu, H., Venugopalan, S., Ramanishka, V., Rohrbach, M., Saenko, K., 2015. A multi-scale multiple instance video description network. arXiv preprint arXiv:1505.05914.

[103] Xu, J., Mei, T., Yao, T., Rui, Y., 2016. Msr-vtt: A large video description dataset for bridging video and language. In: CVPR.

[104] Xu, R., Xiong, C., Chen, W., Corso, J. J., 2015. Jointly modeling deep video and compositional text to bridge vision and language in a unified framework. In: AAAI.

[105] Xu, Z., Yang, Y., Hauptmann, A. G., 2015. A discriminative cnn video representation for event detection. In: CVPR.

[106] Yao, L., Torabi, A., Cho, K., Ballas, N., Pal, C., Larochelle, H., Courville, A., 2015. Describing videos by exploiting temporal structure. In: ICCV.

[107] Yao, T., Mei, T., Ngo, C.-W., 2015. Learning query and image similarities with ranking canonical correlation analysis. In: ICCV.

[108] Yao, T., Mei, T., Ngo, C.-W., Li, S., 2013. Annotation for free: Video tagging by mining user search behavior. In: ACM Multimedia.

[109] Yao, T., Ngo, C.-W., Zhu, S., 2012. Predicting domain adaptivity: Redo or recycle? In: ACM Multimedia.

[110] Yao, T., Pan, Y., Ngo, C.-W., Li, H., Mei, T., 2015. Semi-supervised domain adaptation with subspace learning for visual recognition. In: CVPR.

[111] Ye, G., Li, Y., Xu, H., Liu, D., Chang, S.-F., 2015. Eventnet: A large scale structured concept library for complex event detection in video. In: ACM Multimedia.

[112] Ye, H., Wu, Z., Zhao, R.-W., Wang, X., Jiang, Y.-G., Xue, X., 2015. Evaluating two-stream cnn for video classification. In: ACM ICMR.

[113] Yu, H., Wang, J., Huang, Z., Yang, Y., Xu, W., 2016. Video paragraph captioning using hierarchical recurrent neural networks. In: CVPR.

[114] Zha, S., Luisier, F., Andrews, W., Srivastava, N., Salakhutdinov, R., 2015. Exploiting image-trained cnn architectures for unconstrained video classification. In: BMVC.

[115] Zhang, B., Wang, L., Wang, Z., Qiao, Y., Wang, H., 2016. Real-time action recognition with enhanced motion vector cnns. In: CVPR.

[116] Zhu, W., Hu, J., Sun, G., Cao, X., Qiao, Y., 2016. A key volume mining deep framework for action recognition. In: CVPR.